## REMARKS

Claims 1-14 are currently pending in the above-identified patent application. Claims 1, 7, 11, 13, and 14 have been amended. Claims 1, 7 and 11 have been amended to recite that the storage of data directed to the logical volume during the period when a selected data storage unit is placed in an off-line condition is automatic, while claims 13 and 14 have been amended to correct an obvious typographical error: "apparatus" has been substituted for "method." No new matter has been added by these changes, since support for the changes made to claims 1, 7 and 11 may be found on page 5, lines 20-28, of the subject Specification, as originally filed.

In the subject Office Action, the Examiner rejected claims 1, 7 and 11 under 35 U.S.C. 102(b) as being anticipated by Lubbers et al. (U.S. Patent No. 5,774,643), since the Examiner stated that Lubbers et al. teaches that writes to a data storage unit comprising a logical volume within a RAID storage system may be cached such that in advent of a failure and/or functional cessation of said storage unit (such as may exist upon updating its firmware as claimed), whereby said writes may be subsequently completed upon its corresponding remedy (see Abstract, Figs. 3-8, and Col. 4(5) lines 36(6)); where in view of that it is known to those of ordinary skill in the art and as further reviewed by the reference (see Background of The Invention) that RAID storage systems may be configured to inherently remain fully functional in the presence of the operational cessation of a single redundant element comprising such a system, and/or partially functional upon the failure of otherwise non-further redundant storage elements (i.e. being inherent, in that writes may be cached to an otherwise non-operational storage unit as taught) and thereby may remain on-line until otherwise desired; all claim elements are considered taught and/or inherent in the reference cited although in other form.

Applicants respectfully disagree with the Examiner concerning the rejection of claims 1, 7 and 11 under 35 U.S.C. 102(b) as being anticipated by Lubbers et al. (U.S. Patent No. 5,774,643), for the reasons to be set forth hereinbelow.

Claims 2-6, 8-10 and 12-14 were rejected under 35 U.S.C. 103(a) as being unpatentable over Lubbers et al. (5,774,643), since the Examiner stated that claims

2-6, 8-10, and 12-14, are dependent on claims 1, 7, and 11, respectively, and that Lubbers et al. may be considered to be applicable to any operational cessation of a storage unit as may result from the updating of firmware, diagnostic tests, etc. per claim 2 and applicable to any RAID configuration per claim 3; that said volume may comprise a plurality of disks per claim 4; and that it is implicit that said write cache is logically distinct from said volume per claim 6; that where although Lubbers et al. does not teach that the non-volatile storage utilized to cache writes may comprise a disk drive per claim 5, as disk drives are well known by those of ordinary skill in the art to be non-volatile storage device, the Examiner concluded that it is obvious that the non-volatile write cache as taught by Lubbers et al. may comprise a disk drive for the benefit of cost effectively caching large writes. The Examiner continued that since claims 8-10 and 12-14 are considered encompassed by claims 2-6 in other form, all claims are rejected based upon the same arguments as presented above.

Applicants respectfully disagree with the Examiner concerning the rejection of dependent claims 2-6, 8-10 and 12-14 under 35 U.S.C. 103(a) as being unpatentable over Lubbers et al., U.S. Patent No. 5,774,643, since for the reasons to be set forth hereinbelow, applicants believe that independent claims 1, 7 and 11, from which these claims depend, respectively, are patentable over Lubbers et al. No further discussion of these dependent claims will therefore be provided.

Turning now to the rejection of claims 1, 7 and 11 under 35 U.S.C. 102(b) as being anticipated by Lubbers et al. (U.S. Patent No. 5,774,643), applicants wish to direct the Examiner's attention beginning on Col. 4, line 57, and ending on Col. 5, line 6, wherein it is stated: "[Hereinafter, crash = an unintentional power or function cessation, which could be from a controller, cache, memory drive, computer system etc. unexpectedly ceasing operation due to power loss, failure etc.] In the instance where a crash occurs during a 'write' to disk, it is possible the 'write' wrote some, but not all, of the data blocks and the parity block. This results in inconsistent parity across the slice. Since a write-back cache which is non-volatile is employed, the data that was to be written is still retained, and by some use of the metadata information which is also saved in non-volatile memory, where the write was intended is known. Thus the write that was occurring when the crash took place

can now be reconstructed and it can be insured that the parity is consistent. In this manner only the blocks affected by the crash may be corrected, and parity for every block does not have to be revealed." (emphasis added by applicants). Thus, the write-back cache of Lubbers et al. is continuously running. By contrast, subject claim 1, as amended, recites in part: "... automatically storing data directed to the logical volume in which the selected data storage unit is disposed in a logging repository disposed outside of the logical volume in which the selected data storage unit is disposed during the period of time; ... ." (emphasis added by applicants), and data is stored outside of the logical volume only when the selected data storage unit is being serviced or repaired. Subject claims 7 and 11 have similar recitations.

Further, in the description of FIG. 3, Lubbers et al. state in Col. 8, lines 7-57: "In a dual-redundant configuration, such as the configuration 20 shown in FIG. 3, a controller that is initializing or reinitializing sends information about the process to the other controller. Controllers send keep alive messages to each other at timed intervals. The cessation of communication by one controller causes a 'failover' to occur once the surviving controller has disabled the other controller. In a dual-redundant configuration, if one controller fails, all attached storage devices continue to be served. This is called 'failover'. Failover occurs as has been previously mentioned, because the controllers in a dual-redundant configuration share SCSI-2 device ports and therefore access to all attached storage devices. If a failover is to be achieved, the surviving controller should not require access to the failed controller. The two way failover communication line **37** is depicted in FIG. 3. StorageWorks™ controllers in a dual redundant configuration have the same configuration information at all times. When configuration information is entered into one controller that controller sends the new information to the other controller. Each controller stores this information in a controller resident nonvolatile memory. If one controller fails, the surviving controller continues to serve the failed controller's devices to host computers, thus obviating shared memory access. The controller resolves any discrepancies by using the newest information. Specific firmware components within a controller can communicate with the other controller to synchronize special events between the hardware on both controllers. Some

examples of these special events are SCSI-2 bus resets, cache state changes, and diagnostic tests. Each controller can sense the presence or absence of its cache to set up cache diagnostics and cache operations and can sense the presence or absence of the other controller's cache for dual-controller setup purposes. The failover of a controller's cache occurs only if write-back caching was in use before the controller failure was detected. In this case, the surviving controller causes the failed controller's cache to write its information into the surviving controller's cache. After this is accomplished, the cache is released and access to the devices involved is permitted. The cache then awaits the failed controller's return to the dual-redundant configuration through reinitialization or replacement. If portions of the controller buffer and cache memories fail, the controller continues normal operation. Hardware error correction in controller memory, coupled with advanced diagnostic firmware, allows the controller to survive dynamic and static memory failures. In fact, the controller will continue to operate even if a cache module fails." (emphasis added by applicants).

In a dual redundant configuration, data redundancy is obtained by storing exact copies on mirrored pairs of drives; thus, Lubbers et al. teaches that a failure in one of the dual systems results in the surviving controller writing the cache of the failed controller onto the surviving controller's cache. Therefore, the Lubbers et al. does not teach writing to a data storage unit outside of the logical volume in which the selected data storage unit is disposed during the period of time when the selected data storage unit is being serviced or repaired, as taught by subject claims 1, 7 and 11, as amended. All of the data storage units of Lubbers et al. are within the logical volume of the failed storage unit.

Thus, applicants respectfully believe that Lubbers et al. does not anticipate subject independent claims 1, 7 and 11.

In view of the discussion presented hereinabove, applicants believe that subject claims 1-14, as amended, are in condition for allowance and such action by the Examiner at an early date is earnestly solicited.

Reexamination and reconsideration are respectfully requested.

Respectfully submitted,

COCHRAN FREUND & YOUNG LLC

Date: August 7, 2006

By: _____

Samuel M. Freund
Reg. No. 30,459
2026 Caribou Drive, Suite 201
Fort Collins, Colorado 80525
Phone: (970) 492-1100
Fax: (970) 492-1101
Customer No.: 27479